arXiv:1511.04484v1 [cs.NE] 14 Nov 2015

Unsupervised Learning in Synaptic Sampling Machines

Emre Neftci[1], Bruno Pedroni[2], Siddharth Joshi[3], Maruan Al-Shedivat[4], and Gert Cauwenberghs[2]

[1]Department of Cognitive Sciences, UC Irvine
[2]Department of Bioengineering, UC San Diego
[3]Electrical and Computer Engineering Department, UC San Diego
[4]Machine Learning Department, Carnegie Mellon University

November 17, 2015

Abstract

Recent studies have shown that synaptic unreliability is a robust and sufficient mechanism for inducing the stochasticity observed in cortex. Here, we introduce the Synaptic Sampling Machine (SSM), a stochastic neural network model that uses synaptic unreliability as a means to stochasticity for sampling. Synaptic unreliability plays the dual role of an efficient mechanism for sampling in neuromorphic hardware, and a regularizer during learning akin to DropConnect. Similar to the original formulation of Boltzmann machines, the SSM can be viewed as a stochastic counterpart of Hopfield networks, but where stochasticity is induced by a random mask over the connections. The SSM is trained to learn generative models with a synaptic plasticity rule implementing an event-driven form of contrastive divergence. We demonstrate this by learning a model of MNIST hand-written digit dataset, and by testing it in recognition and inference tasks. We find that SSMs outperform restricted Boltzmann machines (4.4% error rate *vs.* 5%), they are more robust to overfitting, and tend to learn sparser representations. SSMs are remarkably robust to weight pruning: removal of more than 80% of the weakest connections followed by cursory re-learning causes only a negligible performance loss on the MNIST task (4.8% error rate). These results show that SSMs offer substantial improvements in terms of performance, power and complexity over existing methods for unsupervised learning in spiking neural networks, and are thus promising models for machine learning in neuromorphic execution platforms.

1 Introduction

The brain's cognitive power emerges from a collective form of computation extending over very large ensembles of sluggish, imprecise, and unreliable components. This realization spurred scientists and engineers to explore the remarkable mechanisms underlying biological cognitive computing by reverse engineering the brain in "neuromorphic" silicon, providing a means to validate hypotheses on neural structure and function through "analysis by synthesis". Successful realizations of large-scale neuromorphic hardware [43, 58, 33, 50, 65, 11] are confronted by challenges in configuring and deploying them for solving practical tasks, as well as identifying the domains of applications where such devices could excel. The central challenge is to devise a neural computational substrate that can efficiently perform the necessary inference and learning operations in a scalable manner using limited memory resources and limited computational primitives, while relying on temporally and spatially local information.

Recently, one promising approach to configure neuromorphic systems for practical tasks is to take inspiration from machine learning, namely artificial (deep) neural networks. Despite the fact that neural networks were first inspired by biological neurons, the mapping of modern machine learning techniques onto neuromorphic architectures requires overcoming several conceptual barriers. This is because machine learning algorithms and artificial neural networks are designed to operate efficiently on digital processors, often using batch, discrete-time, iterative updates (lacking temporal locality), shared parameters and states (lacking spatial locality), and minimal interprocess communication that optimize the capabilities of GPUs or multicore processors.

Here, we report the SSM, a class of stochastic neural networks that overcomes this conceptual barrier while offering substantial improvements in terms of performance, power and complexity over existing methods for unsupervised learning in spiking neural networks. Similarly to how Boltzmann machines were first proposed [39], the SSM can be viewed as a stochastic counterpart of Hopfield networks [41], but where stochasticity is caused by multiplicative noise.

As in neural sampling [13, 29], neural states (such as instantaneous firing rates or individual spikes) in the SSM are interpreted as Monte Carlo samples of a probability distribution. The source of the stochasticity in neural samplers is often unspecified, or is assumed to be caused by independent, background Poisson activities [60, 56]. Background Poisson activity can be generated self-consistently in balanced excitatory-inhibitory networks [68], or by using finite-size effects and neural mismatch [8]. Although a large enough neural sampling network could generate its own stochasticity in this fashion, the variability in the spike trains decreases strongly as the firing rate increases [30, 52, 17], unless there is an external source of noise or some parameters are fine-tuned. Furthermore, when a neural sampling network generates its own noise, correlations in the background activity can play an adverse role in its performance [62]. Correlations can be mitigated by adding feed-forward connectivity between a balanced network (or other dedicated source of Poisson spike trains) and the neural sampler, but such solutions entail increased resources in neurons, connectivity and ultimately energy. Therefore, the above techniques for generating stochasticity are not ideal for neural sampling.

On the other hand, synaptic unreliability can induce the necessary stochasticity without requiring a dedicated source of Poisson spike trains. The unreliable transmission of synaptic vesicles in biological neurons is a well studied phenomenon [46, 16], and many studies suggested it as a major source of stochasticity in the brain [28, 1, 72, 52]. More recent work suggested synaptic sampling as a mechanism for representing uncertainty in the brain [3], and its role in the synaptic plasticity [45].

We show that uncertainty at the synapse is sufficient in inducing the variability necessary for probabilistic inference in brain-like circuits. Strikingly, we find that the simplest model of synaptic unreliability, called *blank-out* synapses, can vastly improve the performance of spiking neural networks in practical machine learning tasks over existing solutions, while being extremely easy to implement in hardware [34]. In blank-out synapses, for each pre-synaptic spike-event, a synapse evokes a response at the post-synaptic neuron with a probability smaller than one. In the theory of stochastic processes, the operation of removing events from a point process with a probability p is termed p-thinning. Thinning a point process has the interesting property of making the process more Poisson-like. Along these lines, a recent study showed that spiking neural networks endowed with stochastic synapses robustly generate Poisson-like variability over a wide range of firing rates [52], a property observed in the cortex. The iterative process of adding spikes through neuronal integration and removing them through probabilistic synapses causes the spike statistics to be Poisson-like over a large range of firing rates. Other types of noise such as white noise injected to the soma, spike jitter and random synaptic delays do not reproduce such variability.

Synaptic noise in the SSM is not only an efficient mechanism for implementing stochasticity in

spiking neural networks but also plays the role of a regularizer during learning, akin to DropConnect [69]. This technique was used to train artificial neural networks used multiplicative noise on one layer of a feed-forward network for regularization and decorrelation.

Compared to the neural sampler used in [56], the SSM comes at the cost of a quantitative link between the parameters of the probability distribution and those of the spiking neural network. In a machine learning task, this does not pose a problem since the parameters of the spiking neural network can be trained with a learning rule such as event-driven Contrastive Divergence (eCD). ECD is an on-line training algorithm for a subset of stochastic spiking neural networks [56]: The stochastic neural network produces samples from a probability distribution, and Spike Timimg Dependent Plasticity (STDP) carries out the weight updates according to the Contrastive Divergence rule [40] in an online, asynchronous fashion.

SSMs trained with eCD significantly outperform the Restricted Boltzmann Machine (RBM), reaching error rates in an MNIST hand-written digit recognition task of 4.4%, while using a fraction of the number of synaptic operations during learning and inference compared to our previous implementation [56]. Furthermore, the system has two appealing properties: (1) The activity in the hidden layer is very sparse, with less than 10% of hidden neurons being active at any time; (2) The SSM is remarkably robust to the pruning and the down-sampling of the weights: Upon pruning the top 80% of the connections (sorted by weight values) and re-training (32k samples), we observed that the error rate remained below 5%. Overall, the SSM outperforms existing models for unsupervised learning in spiking neural networks while using a fraction of the resources, and thus an excellent candidate for implementation in neuromorphic hardware.

The article is structured as follows: In the Methods section, we first describe the SSM within a discrete-time framework to gain insight into the functionality of the sampling process. We then describe the continuous-time, spiking version of the SSM. In the Results section, we demonstrate the ability of discrete and continuous-time SSMs to learn generative models of the MNIST dataset, and illustrate some of its remarkable features.

2 Materials and Methods

2.1 Discrete-time SSMs as Hopfield Networks with Multiplicative Noise

We first define the discrete-time SSM (dSSM) as a modification of the original Boltzmann Machine. Boltzmann machines are Hopfield networks with thermal noise added to the neurons in order to avoid overfitting and falling into spurious local minima [39]. As in Boltzmann Machines, the dSSM introduces a stochastic component to Hopfield networks. But rather than units themselves being noisy, in the dSSM the connections are noisy.

In the Hopfield network, neurons are threshold units:

$$z_i[t] = \Theta(u_i[t]) = \begin{cases} r_{\text{on}} & \text{if } u_i[t] \geq 0 \\ r_{\text{off}} & \text{if } u_i[t] < 0 \end{cases} \forall i, \tag{1}$$

and connected recurrently through symmetric weights $w_{ij} = w_{ji}$, where $r_{\text{on}} > r_{\text{off}}$ are two numbers indicating the values of the on and off states, respectively.

For simplicity, we chose blank-out synapses as the model of multiplicative noise. With blank-out synapses, the synaptic input to Hopfield unit i is:

$$u_i[t+1] = \sum_{j=1}^{N} \xi_{ij}^{p}[t] z_j[t] w_{ij} + b_i, \tag{2}$$

where $\xi_{ij}^p[t] \in \{0,1\}$ is a Bernoulli process with probability p and b_i is the bias. This corresponds to a weighted sum of Bernoulli variables. Since the $\xi_{ij}^p[t]$ are independent, for large N we can approximate this sum with a Gaussian with mean $\mu_i[t] = b_i + \sum_j pw_{ij}z_j[t]$ and variance $\sigma_i^2[t] = p(1-p)\sum_j (w_{ij}z_j[t])^2$. So,

$$u_i[t+1] = \mu_i[t] + \sigma_i[t]\eta_i[t], \tag{3}$$

where $\eta \sim N(0,1)$. In other terms $u_i[t+1] \sim N(\mu_i[t], \sigma_i^2[t])$. Since $P(z_i[t+1] = 1|\mathbf{z}[t]) = P(u_i[t+1] \geq 0|\mathbf{z}[t])$, the probability that unit i is active given the network state is equal to one minus the cumulative distribution function of u_i:

$$P(u_i[t+1] \geq 0|\mathbf{z}[t]) = \frac{1}{2}\left(1 + \mathrm{erf}\left(\frac{\mu_i[t]}{\sigma_i[t]\sqrt{2}}\right)\right),$$

where "erf" stands for the error function. Plugging back to Eq. (1) gives:

$$\begin{aligned} P(z_i[t+1] = 1|\mathbf{z}[t]) &= P(u_i[t+1] \geq 0|\mathbf{z}[t]) \\ &= \frac{1}{2}\left(1 + \mathrm{erf}\left(\beta\frac{\sum_j w_{ij}z_j[t] + \frac{b_i}{p}}{\sqrt{\sum_j w_{ij}^2 z_j[t]^2}}\right)\right), \\ \beta &= \sqrt{\frac{p}{2(1-p)}}. \end{aligned} \tag{4}$$

Eq. (4) has three interesting properties: 1) the synaptic contributions to each neuron i are normalized such that Euclidian norm of the vector $(w_{i1}z_1, \ldots, w_{iN}z_N)^\top$ is 1; 2) β, which depends on the probability of the synaptic blank-out noise and the network state, plays the role of thermal noise in the Boltzmann machine; and 3) In general, despite that the connectivity matrix is symmetric, the interactions in the SSM are asymmetric because the effective slope of the sigmoid function depends on the normalizing factor of each neuron.

In the general case, the three properties above imply that the SSM cannot be expressed in terms of an energy-based model (as in the case for Boltzmann machines), and therefore we cannot derive the joint distribution $P(\mathbf{z}[\mathbf{t}])$. Some insight can be gained in the particular case where $r_{\mathrm{on}} = -r_{\mathrm{off}}$. In this case, due to the square exponent of the neural state $z_i[t]$, the denominator in Eq. (4) simplifies such that:

$$P(z_i[t+1] = 1|\mathbf{z}[t]) = P(u_i[t+1] \geq 0|\mathbf{z}[t]) \tag{5}$$

$$= \frac{1}{2}\left(1 + \mathrm{erf}\left(\beta\frac{\sum_j w_{ij}z_j[t] + \frac{b_i}{p}}{r_{\mathrm{on}}\sqrt{\sum_j w_{ij}^2}}\right)\right). \tag{6}$$

In other words: the variability in the network does not depend on the neural states. With the additional constraint that $\sum_i w_{ij}^2 = \sum_j w_{ij}^2$, the connectivity matrix becomes symmetric. In this case, the resulting system is *almost* a Boltzmann machine, with the only exception that the neural activation function is an error function instead of the logistic function. Generally speaking, the error function is a sigmoid function which is reasonably close to the logistic function after a rescaling of the argument. Therefore the parameters of a Boltzmann machine can be approximately mapped on the discrete-time SSM with $r_{\mathrm{on}} = -r_{\mathrm{off}}$. To test the quality of this approximation, we compute the Kullback-Leibler (KL) divergence between an exact Boltzmann distribution and the distribution samples in the case of a discrete-time SSM with $r_{\mathrm{on}} = 1, r_{\mathrm{off}} = -1$. As expected, the KL-divergence

between the dSSM and the exact distribution reaches a plateau due to the approximation caused by the inexact error function. However, the results show that in networks of this size the distribution sampled by the dSSM has the same KL-divergence as the RBM obtained after 10^5 iterations of Gibbs sampling, which is more than the typical number of iterations used for many tasks [37].

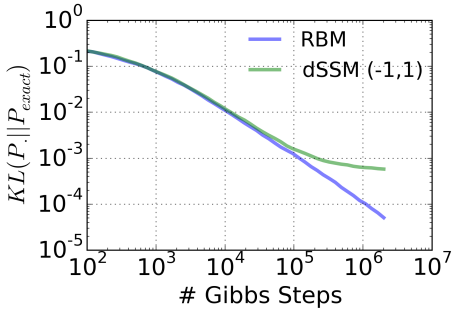

Figure 1: KL-divergence between a restricted Boltzmann distribution (P_{exact}, 5 hidden units, 5 visible units) and the distribution sampled by a discrete-time SSM with $r_{\text{on}} = 1, r_{\text{off}} = -1$ and matched parameters. As a reference, the KL-divergence using an RBM is shown (blue curve). The saturation of the green curve is caused by the inexact activation function of the SSM unit (error function instead of the logistic function).

We premise that the SSMs with all or none activation values ($r_{\text{on}} \neq r_{\text{off}} = 0$) behave in a manner that is sufficiently similar, so that learning in the dSSM with CD is approximately valid. In the Results section, we test this premise on the MNIST hand-written digit machine learning task.

2.1.1 Learning rule for RBM and dSSMs

The training of RBMs proceeds in two phases. At first the states of the visible units are clamped to a given vector of the training set, then the states of the hidden units are sampled. In a second "reconstruction" phase, the network is allowed to run freely. Using the statistics collected during sampling, the weights are updated in a way that they maximize the likelihood of the data. Collecting equilibrium statistics over the data distribution in the reconstruction phase is often computationally prohibitive. The Contrastive Divergence (CD) algorithm has been proposed to mitigate this [40, 38]: the reconstruction of the visible units' activity is achieved by sampling them conditioned on the values of the hidden units. This procedure can be repeated k times (the rule is then called CD_k), but relatively good convergence is obtained for the equilibrium distribution even for one iteration. The CD learning rule is summarized as follows:

$$\Delta w_{ij} = \epsilon(\langle v_i h_j \rangle_{\text{data}} - \langle v_i h_j \rangle_{\text{recon}}), \tag{7}$$

where v_i and h_j are the activities in the visible and hidden layers, respectively. The learning of the biases follows a similar rule.

The learning rule for the dSSM was identical to that of the RBM.

2.2 Continuous-time, Spiking SSM

Neuron and Synapse Model

Except for the noise model, the neuron and synapse models used in this work are identical to those described in [56] and are summarized here for convenience. The neuron's membrane potential below

firing threshold θ is governed by the following differential equation:

$$C\frac{\mathrm{d}}{\mathrm{d}t}u_i = -g_L u_i + I_i(t), \quad u_i(t) \in (-\infty, \theta), \tag{8}$$

where C is a membrane capacitance, u_i is the membrane potential of neuron i, g_L is a leak conductance, $I_i(t)$ is the time-varying input current and θ is the neuron's firing threshold. When the membrane potential reaches θ, an action potential is elicited. After a spike is generated, the membrane potential is clamped to the reset potential u_{rst} for a refractory period τ_r.

In the SSM, the currents $I_i(t)$ depend on the layer the neuron is situated in (visible v or hidden h). For a neuron i in the visible layer v

$$\begin{aligned} I_i(t) =& I_i^d(t) + I_i^v(t), \\ \tau_{syn}\frac{\mathrm{d}}{\mathrm{d}t}I_i^v(t) =& -I_i^v + \sum_{j=1}^{N_h} \xi_{h,ji}^p(t) q_{ji} h_j(t) + b_{v_i}, \end{aligned} \tag{9}$$

where $I_i^d(t)$ is a current representing the data (*i.e.* the external input), $I^v(t)$ is the feedback from the hidden layer activity and the bias. The q's are the respective synaptic weights, $\xi_{v,ij}^p$ is a Bernoulli process (*i.e.* "coin flips") with probability p implementing the stochasticity at the synapse, and b_{v_i} are constant currents representing the biases of the visible neurons. Spike trains are represented by a sum of Dirac delta pulses centered on the respective spike times:

$$h_j(t) = \sum_{k \in Sp_j^h} \delta(t - t_k), \quad v_i(t) = \sum_{k \in Sp_i^v} \delta(t - t_k) \tag{10}$$

where Sp_j^h, Sp_j^v are the sets the spike times of the hidden neuron h_j and visible neurons v_i, respectively, and $\delta(t) = 1$ if $t = 0$ and 0 otherwise.

For a neuron j in the hidden layer h,

$$\begin{aligned} I_j(t) =& I_j^h(t), \\ \tau_{syn}\frac{\mathrm{d}}{\mathrm{d}t}I_j^h(t) =& -I_j^h + \sum_{i=1}^{N_v} \xi_{v,ij}^p(t) q_{ij} v_i(t) + b_{h_j}, \end{aligned} \tag{11}$$

where $I^h(t)$ is the feedback from the visible layer, and $v(t)$ are Poisson spike trains of the visible neurons defined in Eq. (10) and b_{h_j} are the biases of the hidden neurons. The dynamics of I^h and I^v correspond to a first-order linear filter, so each incoming spike results in Post–Synaptic Potentials (PSPs) that rise and decay exponentially (*i.e.* alpha-PSP) [32].

2.2.1 Event-Driven CD Synaptic Plasticity Rule

Event-driven Contrastive Divergence is an online variation of CD amenable for implementation in neuromorphic hardware: by interpreting spikes as samples of a probability distribution, a possible neural mechanism for implementing CD is to use STDP synapses to carry out CD-like updates.

The weight update in eCD is a modulated, pair-based STDP rule with learning rate ϵ_q:

$$\frac{\mathrm{d}}{\mathrm{d}t}q_{ij} = \epsilon_q g(t)\, \mathrm{STDP}_{ij}(v_i(t), h_j(t)) \tag{12}$$

where $g(t) \in \mathbb{R}$ is a zero-mean global gating signal controlling the data vs. reconstruction phase, q_{ij} is the weight of the synapse and $v_i(t)$ and $h_j(t)$ refer to the spike trains of neurons v_i and h_j, defined as in Eq. (10). The same rule is applied to learn biases $\mathbf{b}_v$ and $\mathbf{b}_h$:

$$\frac{\mathrm{d}}{\mathrm{d}t} b_{v_i} = \epsilon_b \frac{1}{2} g(t)\, \mathrm{STDP}_i(v_i(t), v_i(t)), \tag{13}$$

$$\frac{\mathrm{d}}{\mathrm{d}t} b_{h_i} = \epsilon_b \frac{1}{2} g(t)\, \mathrm{STDP}_i(h_i(t), h_i(t)), \tag{14}$$

where ϵ_b is the learning rate of the biases. The weight update is governed by a symmetric STDP rule with temporal window $K(t) = K(-t), \forall t$:

$$\begin{aligned} \mathrm{STDP}_{ij}(v_i(t), h_j(t)) =& v_i(t) A_{h_j}(t) + h_j(t) A_{v_i}(t), \\ A_{h_j}(t) =& A \int_{-\infty}^{t} \mathrm{d}s K(t-s) h_j(s), \\ A_{v_i}(t) =& A \int_{-\infty}^{t} \mathrm{d}s K(s-t) v_i(s), \end{aligned} \tag{15}$$

with $A > 0$ defining the magnitude of the weight updates.

Although any symmetric learning window can be used, for simplicity, we used a nearest neighbor update rule where:

$$K(t-s) = \begin{cases} 1 & \text{if } |t-s| < \tau_{STDP} \\ 0 & \text{otherwise} \end{cases}$$

In our implementation, updates are additive and weights can change polarity. A global modulatory signal that is synchronized with the data clamping phase modulates the learning to implement CD:

$$g(t) = \begin{cases} 1 & \text{if } mod(t, 2T) \in (\tau_{br}, T) \\ -1 & \text{if } mod(t, 2T) \in (T + \tau_{br}, 2T) \\ 0 & \text{otherwise.} \end{cases}, \tag{16}$$

The signal $g(t)$ switches the behavior of the synapse from Long-Term Potentiation (LTP) to Long-Term Depression (LTD) (*i.e.* Hebbian to Anti-Hebbian). The temporal average of $g(t)$ vanishes to balance LTP and LTD. The modulation factor is zero during some time τ_{br}, so that the network samples closer to its stationary distribution when the weights updates are carried out. The time constant τ_{br} corresponds to a "burn-in" time of MCMC sampling and depends on the overall network dynamics. [56] showed that, when pre-synaptic neurons and post-synaptic neurons fire according to Poisson statistics, eCD is equivalent to CD. The effective learning window is:

$$\epsilon = 2A \frac{T - \tau_{br}}{2T}.$$

We note that the learning rate between the dSSM and the spiking SSM cannot be compared directly because the relation between the synaptic weights q and w is not known.

In the spiking SSM, weight updates are carried out even if a spike is dropped at the synapse. This speeds up learning without adversely affecting the entire learning process because spikes dropped at the synapses are valid samples in the sense of the sampling process. During the data phase, the visible units were driven with constant currents equal to the logit of the pixel intensity (bounded to the range $[10^{-5}, .98]$ in order to avoid infinitely large currents), plus a white noise process of low amplitude σ to simulate sensor noise.

2.2.2 Synaptic Blank-out Noise

Perhaps the simplest model of synaptic noise is the blank-out noise: for each spike-event, the synapse has a probability p of evoking a PSP, and a probability $(1-p)$ that of evoking no response. For a Poisson spike train of rate ν, the synaptic stochasticity gives a Poisson spike train with rate $p\nu$ [34, 59]. For a periodic (regular) spike train, stochastic synapses add stochasticity to the system. The coefficient of variation of the Inter–Spike Interval (ISI) becomes $\sqrt{1-p}$. The periodic spike train thus tends to a Poisson spike train when $p \to 0$, with the caveat that events occur at integer multiples of the original ISI. Intuitively, the neural network cycles through deterministic neural integration and stochastic synapses. [52] found that the recurring process of adding spikes through neuronal integration and removing them through stochastic synapses causes the spike statistics to remain Poisson-like (constant Fano Factor) over a large dynamical range. Synaptic stochasticity is thus a biologically plausible source of stochasticity in spiking neural networks, and can be very simply implemented in software and hardware [34].

2.2.3 Spiking neural networks with Poisson input and blank-out synapses

This section describes the neural activation function of leaky Integrate & Fire (I&F) neurons. The collective dynamics of spiking neural circuits driven by Poisson spike trains is often studied in the diffusion approximation [71, 18, 17, 30, 63, 24, 67]. In this approximation, the firing rates of individual neurons are replaced by a common time-dependent population activity variable with the same mean and two-point correlation function as the original variables, corresponding here to a Gaussian process. The approximation is true when the following assumptions are verified: 1) the charge delivered by each spike to the post-synaptic neuron is small compared to the charge necessary to generate an action potential, 2) there is a large number of afferent inputs to each neuron, 3) the spike times are uncorrelated. In the diffusion approximation, only the first two moments of the synaptic current are retained. The currents to the neuron, $I(t)$, can then be decomposed as:

$$I(t) = \mu + \sigma\eta(t), \tag{17}$$

where $\mu = \langle I(t) \rangle = p\sum_j q_j \nu_j + b$ and $\sigma^2 = p\sum_j q_j^2 \nu_j$, ν_j is the firing rate of pre-synaptic neuron j, and $\eta(t)$ is the white noise process. Note that a Poisson spike train of mean rate ν with spikes removed with probability p is a Poisson spike train with parameter $p\nu$. As a result, blank-out synapses have the effect of scaling the mean and the variance by p. The firing rate ν_j is interpreted as a scaled version of the conditional probabilities: $\frac{1}{\tau_{ref}} p(z_j(t) = 1 | \mathbf{z}(t))$.

In this case, the neuron's membrane potential dynamics is an Ornstein-Uhlenbeck process [31]. For simplicity of presentation, the reset voltage was set to zero ($u_{rst} = 0$). We consider the case where the synaptic time constant dominates the membrane time constant. In other words, the membrane potential closely follows the dynamics of the synaptic currents and the effect of the firing threshold and the reset in the distribution of the membrane potential can be neglected. This problem was studied in great detail with first order approximations in τ_m/τ_{syn} [60]. For comparison with the dSSM, we focus here on the case $\tau_m = 0$. In this case, the stationary distribution is a Gaussian distribution:

$$u \sim N(\frac{\mu}{g_L}, \frac{\sigma^2}{2g_L^2 \tau_{syn}}).$$

Neurons such that $p(u > 0|\nu)$ fire at their maximum rate of $\frac{1}{\tau_{ref}}$. Following similar steps as in the

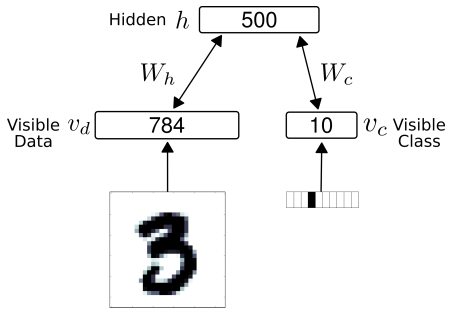

Figure 2: The network architecture consists of a visible and a hidden layer. The visible layer is partitioned into 784 sensory (data) neurons ($\mathbf{v_d}$) and 10 class label neurons ($\mathbf{v_c}$). During data presentation, the activities in the visible layer are clamped to the data, consisting of a digit and its label (one-hot coded). In the RBM and the SSM, the weight matrix between the visible layer and the hidden layer is symmetric.

dSSM, the firing rate of a neuron i becomes:

$$\begin{aligned}\nu_i &= \frac{1}{\tau_{ref}}\frac{1}{2}\left(1+\mathrm{erf}\left(\frac{\mu_i}{\sigma_i}\sqrt{\tau_{syn}}\right)\right), \\ &= \frac{1}{\tau_{ref}}\frac{1}{2}\left(1+\mathrm{erf}\left(\sqrt{p}\frac{\sum_j q_{ij}\nu_j+\frac{b_i}{p}}{\sqrt{\sum_j q_{ij}^2\nu_j}}\sqrt{\tau_{syn}}\right)\right).\end{aligned} \tag{18}$$

Eq. (18) clarifies how the noise amplitude σ affects the neural activation ν, and thus allows a quantitative comparison with the dSSM. Similarly to Eq. (4), the input is effectively normalized by the variability in the inputs (where on/off values are replaced by firing rates). These strong similarities suggest that the dSSM and the spiking neural network are equivalent. Using computer simulations of the MNIST learning task, in the Results section we show that the two networks indeed perform similarly.

2.3 Training protocol and network structure for MNIST

RBM and the discrete-time SSM The network consisted of 1294 neurons, partitioned into 794 visible neurons and 500 hidden neurons (Fig. 2). The visible neurons were partitioned into 784 sensory (data) neurons and 10 class label neurons. The discrete-time SSM (dSSM) is similar to the RBM in all manners, except that the units are threshold functions, and each weight is multiplied by independently drawn binomials ($p = .5$) at every step of the Gibbs sampling.

In both cases the training set was randomly partitioned into batches of equal size N_{batch}. One epoch is defined as the presentation of the full MNIST dataset (60000 digits). We used CD-1 to train the RBM and the dSSM.

Classification is performed by choosing the most likely label given the input, under the learned model. To estimate the discrimination performance, we sampled the distribution of class units $P(class|digit)$ using multiple Gibbs Sampling chains (50 parallel chains, 2 steps). We take the classification result to be the label associated to the most active class unit averaged across all chains.

For testing the discrimination performance of an energy-based model such as the RBM, it is

common to compute the free-energy $F(\mathbf{v_c})$ of the class units [36], defined as:

$$\exp(-F(\mathbf{v_c})) = \sum_{\mathbf{v_d},\mathbf{h}} \exp(-E(\mathbf{v_d}, \mathbf{v_c}, \mathbf{h})), \tag{19}$$

and selecting $\mathbf{v_c}$ such that the free-energy is minimized. The free-energy computes the probabilities of a given neuron to be active, using the joint distribution of the RBM. Classification using free energy is inapplicable to SSMs because it cannot be expressed in terms of an energy-based model (See Methods). Therefore, throughout the article, free energy-based classification was used only for the RBM.

The learning rate in the RBM and the dSSM was linearly reduced during the training, reaching 0 at the end of the learning. Both RBM and dSSM were implemented on a GPU using the Theano package [12].

Continuous-time SSM The spiking SSM network structure is identical to that of the dSSM. Similarly to [56], for a given digit, each neuron in layer $\mathbf{v}$ was injected with currents transformed according to a logit function $I_i^d \propto \log\left(\frac{s_i}{1-s_i\tau_r}\right)$, where s_i is the value of pixel i. To avoid saturation of the neurons using the logit function, the pixel values of the digits were bounded to the range $[10^{-5}, .98]$

Training followed the eCD rule described above. The learning rate was decayed linearly during the training, reaching 0 at the end of the learning. Similarly to the dSSM, the discrimination performance is evaluated by sampling the activities of the class units for every digit in the test dataset. In the spiking SSM, this is equivalent to identifying the neuron that has the highest firing rate. The spiking neural network was implemented in the spike-based neural simulator Auryn, optimized for recurrent spiking neural networks with synaptic plasticity [73]. Connections in the SSM are symmetric. But due to the constraints in the parallel simulator, the connections were implemented using two separate synapses (one in each direction), and periodically symmetrized to maintain symmetry (every 1000 s of simulated time). A complete list of parameters used in our simulator is provided in the appendix Tab. 2.

3 Results

The Boltzmann machine is a neural network consisting of stochastic binary neurons, where the probability distribution over the binary states is the Boltzmann distribution. The Boltzmann machine was introduced as a stochastic variation of the Hopfield network [39]. Stochasticity caused by thermal noise would randomly flip the state of a neuron to perform more robustly by avoiding local minima of the energy function. The SSM is a similar extension of the Hopfield network. The key idea is to use multiplicative noise at the connections (synapses) as a source of stochasticity. We show in the Methods section that multiplicative noise has the property of dynamically modifying the slope of the activation function of the neurons.

We demonstrate two different implementations of the SSM, one continuous-time (spiking) and one discrete-time (non-spiking, Hopfield network-like). The discrete-time SSM (dSSM) is a variant of the original RBM, used to provide insight into the functionality of the SSM, and to justify the continuous-time, spiking SSM. As in RBMs, the dSSM was sampled using Gibbs sampling. The continuous-time, spiking SSM consisted of a network of deterministic I&F spiking neurons, connected through stochastic (blank-out) synapses. As in [56], spikes in the network are interpreted as random vectors (samples) of an underlying probability distribution. The RBM and the dSSM were trained using the Contrastive Divergence (CD) learning rule, and the spike-based SSM was trained

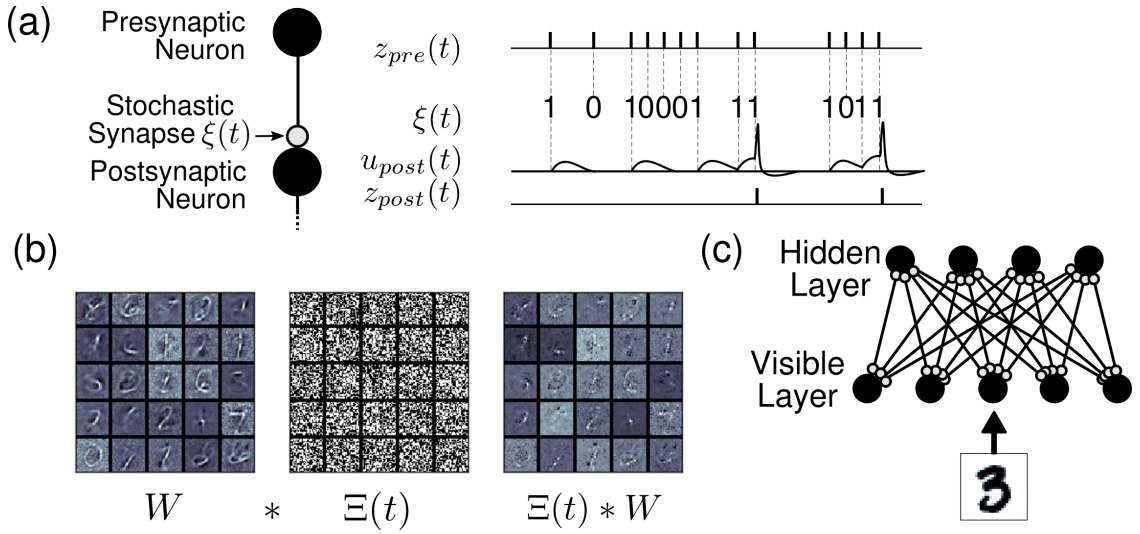

Figure 3: Illustration of the Synaptic Sampling Machine (SSM). (a) At every occurrence of a pre-synaptic event, a pre-synaptic event is propagated to the post-synaptic neuron with probability p. (b) Synaptic stochasticity can be viewed as a continuous DropConnect method where weights are masked by a binary matrix M(t), where $*$ denotes element-wise multiplication. (c) SSM Network architecture, consisting of a visible and a hidden layer. For classification experiments, the visible layer can be partitioned in a data part and a label part.

using event-driven Contrastive Divergence. Contrary to the model presented in [56], the probability distribution sampled by the SSM is not known. This is because the variability introduced by blank-out synapses depends on the network states, which significantly complicates the mathematical formulation of the sampling procedure (Sec. 2.2.3). In the Methods section, we show that discrete-time SSMs with activation values $(-1, +1)$ are similar to Boltzmann machines, except that the activation function of the neuron is an error function instead of the logistic function, and where the input is invariant to rescaling. We assumed that the dSSM with activation levels (0,1) and the spiking SSM are sufficiently similar to the Boltzmann distribution, such that CD and eCD as originally described in [56] are approximately valid. We test this assumption through computer simulations of the SSMs in an MNIST hand-written digit learning task.

3.1 Unsupervised Learning of MNIST handwritten digits in Synaptic Sampling Machines

Similarly to RBMs, SSMs can be trained as generative models. Generative models have the advantage that they can act simultaneously as classifiers, content-addressable memories, and carry out probabilistic inference. We demonstrate these features in a MNIST hand-written digit task [47], using networks consisting of two layers, one visible and one hidden (Fig. 2).

Learning results are summarized in Tab. 1. The spiking SSM appears to slightly outperform the discrete-time SSM, although a direct comparison between the two is not possible because the sampling mechanism and the batch sizes are different. We find that SSMs attain classification accuracies that slightly outperform the machine learning algorithm (error rates: spiking SSM 4.4% *vs.* RBM 5%), even after much fewer repetitions of the dataset (Fig. 4). To date, this is the best performing spike-based unsupervised learner on MNIST. The spiking network implementing the SSM is many times smaller than the best performing spike-based unsupervised learner to date

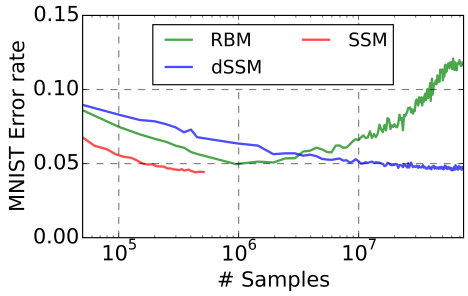

Figure 4: *The SSM outperforms its RBM counterpart at the MNIST task.* The RBM, the discrete-time SSM (dSSM) and the continuous-time (spiking) SSM were trained to learn a generative model of MNIST handwritten digits. The discrete-time SSM model is identical to the RBM except that it consists of threshold units with stochastic blank-out synapses with probability 50%. The recognition performance is assessed on the testing dataset which was not used during training (10000 digits). Error rate in the RBM starts increasing after reaching peak performance, mainly due to decreased ergodicity of the Markov chains and overfitting. Learning in the dSSM is slower than in the RBM, as reported with other models using DropConnect [69] but it is effective in preventing overfitting. At the end of the training, the recognition performance of the spiking SSMs averaged over 8 runs with different seeds reached 4.6% error rate. Due to the computational load of running the spike-based simulations on the digital computer, the spiking SSM was halted earlier than the RBM and the dSSM. In spite of the smaller number of digit presentations, the spiking SSM outperformed the RBM and the dSSM. This is partly because weight updates in the spiking SSM are undertaken during each digit presentation, rather than after each minibatch.

([1294 neurons 800k synapses vs 7184 neurons, 5M synapses]) [25]. For comparisons with other recent techniques for unsupervised learning in spiking neural networks, we refer the reader to [25], which provides a recent survey of the MNIST learning task with spiking neural networks.

We tested the speed of digit classification under the trained SSM. We computed the prediction after sampling for a fixed time window that we varied from 0ms to 300ms (Fig. 5). Results show that the trained network required approximately 250ms to reach the peak performance of 95.8%, and that already 87% of the predictions were correct after 50ms.

Similarly to RBMs, the SSM learns a generative model of the MNIST dataset. This generative model allows to generate digits and reconstruct them when a part of the digit has been occluded. We demonstrate this feature in a pattern completion experiment where the right half of each digit was presented to the network, and the visible neurons associated to the left half of the digit were

Model, $n_H = 500$, 60k digits training set	MNIST Accuracy
Synaptic Sampling + Event-driven CD (this work)	**95.6**%
Synaptic Sampling + Event-driven CD + 81.5% connection pruning + relearning (this work)	**95.4**%
Synaptic Sampling + Event-driven CD + 4 bit synaptic weights (this work)	94.8%
Synaptic Sampling + Event-driven CD + 2 bit synaptic weights (this work)	92.2%
Discrete-time Synaptic Sampling + Standard CD (this work)	**95.5**%
Gibbs Sampling + Standard CD	95.0%
Neural Sampling + Event-driven CD [56]	91.9%
Neural Sampling + Event-driven CD [56] (5 bits post-learning rounding)	89.4%

Table 1

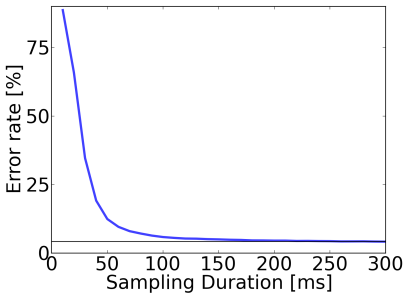

Figure 5: *Accuracy evaluation in the spiking SSM.* To test recognition accuracy, for each digit in the test dataset (10000 digits), class neurons in the SSMs are sampled for up to 300 ms. The classification is read out by identifying the group of class label neurons that had the highest activity and averaging over all digits of the test set. The accuracy after 50 ms of sampling was above 87% and after 250 ms the accuracies typically reached their peak for this trained network (95.8%, horizontal bar).

sampled (Fig. 6). In most cases, the SSM correctly completed the left half of the digit. Fig. 6 also illustrates the dynamics of the completion, which appears to reach a stationary state after about 200ms. Overall, these results show that the SSM can achieve similar tasks as in RBMs, at a similar or better recognition performance while requiring fewer dataset iterations during learning.

3.2 Representations Learned in SSMs are Sparse

In deep belief networks, discriminative performance can improve when using binary features that are only rarely active [54]. Furthermore, in hardware sparser representations result in lower communication overhead, and therefore lower power consumption.

To quantify the degree of sparsity in the RBM and the SSM, we computed the average fraction of active neurons. For a similar parametrization, the average activity ratio of SSMs is much lower than that of RBMs (Fig. 7). RBMs can be trained to be sparse by added sparsity constraints during learning [48], but how such constraints can map to spiking neurons and synaptic plasticity rules is not straightforward. In this regard, it is remarkable that SSMs are sparse without including any sparsity constraint in the model.

One can gain an intuition on the cause for sparsity in the SSM by examining the spiking dynamics during learning: In the absence of additive noise, the input-output profile of leaky I&F neurons near the rheobase is rectified-linear. Consequently, without positive inputs and positive bias values, the spiking neuron cannot fire, and thus the pre-synaptic weights cannot potentiate. By selecting bias values at or below zero, this feature causes neurons in the network to be progressively recruited, thereby promoting sparsity. This is to be contrasted in the case of neurons with additive noise (such as white noise with constant amplitude), which can fire even if the inputs are well below rheobase. In the SSM, the progressive recruitment can be observed in Fig. 8: early in the training, hidden neurons are completely silent during the reconstruction phase. Because there is no external stimulus during the reconstruction phase, only neurons that were active during the data phase are active during the reconstruction phase. After the presentation of 60 digits, the activity appears to "grow" from the data phase.

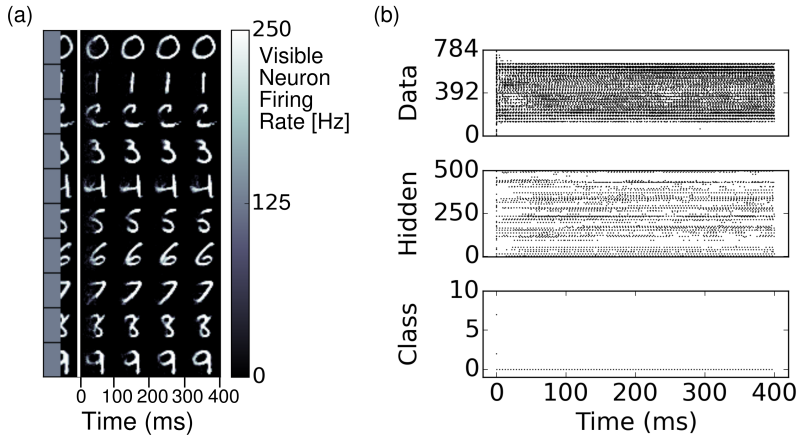

Figure 6: Pattern completion in the spiking SSM. (a) The right half of a digit and its label is presented to the visible layer of the SSM. The ensuing activity in the visible layer is partitioned in 100 ms bins and the average activity during each bin is presented in the color plots to the right. The network completes the left half of the digit. (b) Raster plot of the pattern completion for the first row of panel (a).

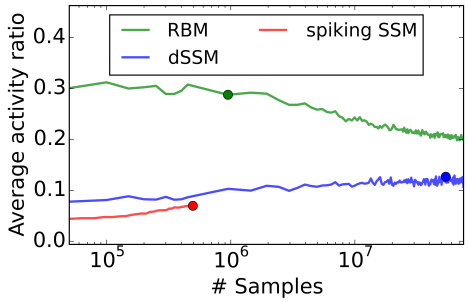

Figure 7: *For a similar parametrization, SSMs learn sparser representations than RBMs.* The curves plot the average number of active units during learning in the hidden layer, computed over 50 digits for the dSSM and the RBM, and 1000 digits for the spiking SSM. In the case of the spiking SSM, the average activity ratio is the average firing rate of the neurons in the network divided by the maximum firing rate t_{ref}^{-1}. The colored dot indicates the point where the peak of the recognition accuracy was reached. Nearly three times fewer units were active in the SSM than in the RBM at the designated points, and even fewer in the case of the spiking SSM.

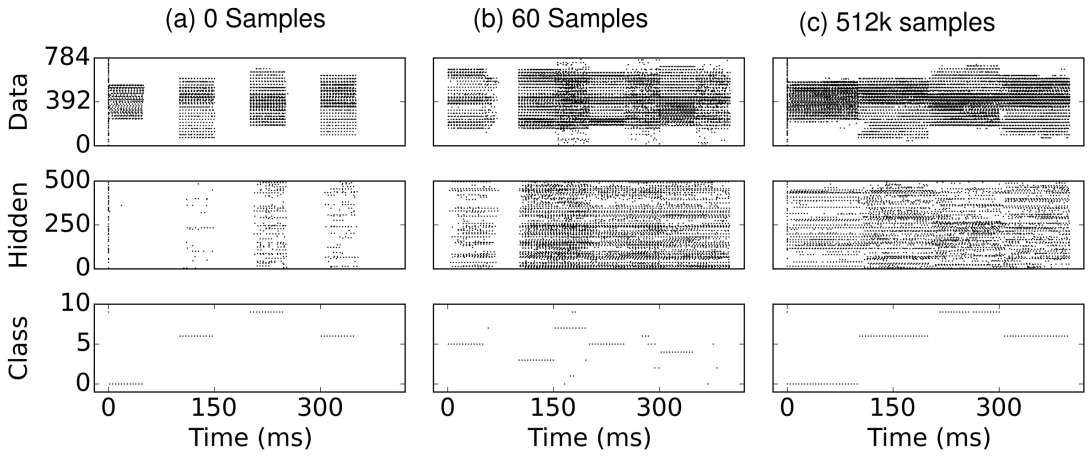

Figure 8: (a-c) Spike rasters during eCD learning. The data phase is $0-50$ ms and the reconstruction phase is $50-100$ ms, repeated every 100 ms for a different digit of the training set. Due to the deterministic neural dynamics, a neuron receiving no input can never fire if its bias is below the rheobase (*i.e.* the minimum amount of current necessary to make a neuron fire). Consequently, at the beginning of learning, the hidden neurons are completely silent in the reconstruction phases, and are gradually recruited during the data phase of the eCD rule (*e.g.* see 50 ms in panel (b)).

3.3 Robustness of the SSM to synapse pruning and weight down-sampling

Sparse networks that require using very few bits for storage can have a lower memory cost, along with a smaller communication and energy footprint. To test the storage requirements for the spiking SSM and its robustness to pruning connections, we removed all synapses whose weights were below a given threshold. We varied the threshold such that the ratio of remaining connections spanned the range $[0\%, 100\%]$.

The resulting accuracy, plotted against the ratio of connections retained, is shown in Fig. 9a (green curve). We then re-trained these networks with SSM over 32 epochs testing once every 4 epochs. The re-trained models were tested against $10,000$ images from the MNIST dataset. The re-learning substantially recovered the performance loss caused by the weight pruning as shown in Fig. 9a (red curve). This result suggests that only a relatively small number of connections play a critical role in the underlying model for the recognition task.

3.3.1 Learning low precision weight synapses

Dedicated memory for storing the synaptic connectivity table and the synaptic weights often occupies the largest area in neuromorphic devices [51, 50]. Therefore, the ability to reduce the synaptic precision required for the operation of an algorithm can be very beneficial in hardware. We quantify the effects of lowered precision in inference by downsampling the synaptic weights, while performing computation at full precision. In the context of a hardware implementation, this results in lower memory costs since fewer bits can be used to store the same network. To test the performance of the asynchronous SSM with lowered resolution, we truncated the weights to a single decimal point. The resulting weights were restricted to less than 128 unique values (7 bits). In Fig. 9b, we truncated the weights of the previously pruned network to 7 bits, and examined the results over the range of retained connections. The results of testing 10,000 samples of MNIST on this network are shown in Fig. 9b. The error rate was about 8% with 12.5% of the total synaptic weights retained and about

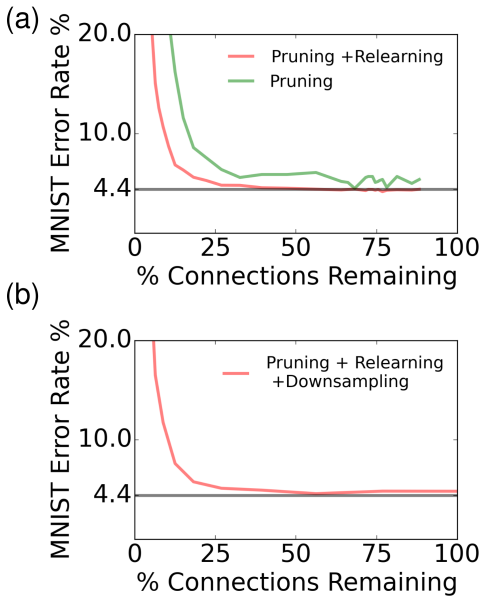

Figure 9: *Pruning connections in the SSM and re-learning.* (a) We test the spiking SSM's robustness to synapse pruning by removing all connections whose weights were below a given threshold. We quantified the effects of pruning by testing the pruned SSM over the entire MNIST test set (green curve). After pruning, we re-trained the network over 32 epochs (corresponding to 32k sample presentations), which recovered most of the loss due to pruning (red curve). (b) Reduced synaptic weight precision in the SSM. After learning, pruning and re-learning, the synaptic weights were downsampled by truncating them to a single decimal point. There remained fewer than 128 unique weight values in the final matrix. The horizontal line is at 4.4%.

5% with 49% of the connections retained.

Another recently proposed technique for reducing the precision of the weights while minimizing impact on performance is the dual-copy rounding technique [66]. In the context of our sampling machine, the key idea is to sample using reduced precision weights, but learn with full precision weights. Dual copy rounding was shown to outperform rounding of the weight after learning.

Training the spiking SSM with dual-copy rounding at 4-bit weights (16 different weight values) resulted in 94.8% accuracy at the MNIST task, and rounding at 2 bits (4 weight values) reduced this number to 92.2%. Synaptic weight resolution of 4 bits is recognized as a sweet spot for hardware [50, 61]. Furthermore, it is a plausible synaptic weight precision for biological synapses: recent analysis of synapses in the rat hippocampus suggested that each synapse could store about 4 to 5 bits of information [10].

We note that the dual-copy approach is only beneficial at the inference stage (post-learning). During inference, the high-precision weights can be dropped and only the low-precision weight are maintained. However, during learning it is necessary to maintain full precision weights. Learning with low precision weights is possible using stochastic rounding techniques [53]. We could not test stochastic rounding in spiking SSMs because the symmetry requirements in the spiking neural network connectivity prevent a direct, efficient implementation in multithread simulators (such as the used Auryn neural simulator).

3.4 Synaptic Operations in the SSMs

Power consumption in neuromorphic computing is often dominated by synaptic communication and plasticity. Akin to the multiply accumulate (MAC), the synaptic operation (SynOp) is thus a representative metric of the performance and energy efficiency of neuromorphic hardware [50].

A SynOp potentially consumes much less energy than a MAC even on targeted GPGPUs, and improvements in energy per SynOP directly translate into energy efficiencies at the level of the task. However, these operations are not directly comparable because it is unclear how many SynOps provide the equivalent of a MAC operation at a given task. To provide a reference to this comparison, we count the number of SynOps and approximate number of MACs necessary to achieve a target accuracy at the MNIST task for SSMs and RBMs, respectively (Fig. 10). Strikingly, the SSM achieves SynOp-MAC parity at this task. Given that the energy efficiency of a SynOp is potentially orders of magnitude lower than the MAC [50], this result makes an extremely strong case for hardware implementations of SSMs. The possible reasons for this parity are twofold: 1) weight updates are undertaken after presentation of every digit, which mean that fewer repetitions of the dataset are necessary to achieve the same performance. 2) only active connections incur a SynOp. In the RBM, the operations necessary for computing the sigmoid function, random numbers at the neuron, and the weight updates were not taken into account and would further favor the spiking SSM. The operation necessary for the stochastic synapse in the SSM (a Bernoulli trial) is likely to be minimal because the downstream synapse circuits do not need to be activated when the synapse fails to transmit. Furthermore, the robustness of SSMs to the pruning of connections (described in Sec. 3.3) can further strongly reduce the number of SynOps after learning.

4 Discussion

The Boltzmann Machine stems from the idea of introducing noise into a Hopfield network [39]. The noise prevents the state from falling into local minima of the energy, enabling it to learn more robust representations while being less prone to overfitting. Following the same spirit, the SSM introduces

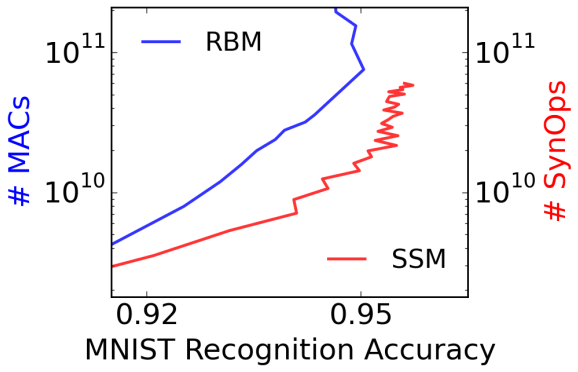

Figure 10: *The SSM achieves SynOp-MAC parity at the MNIST task.* The number of multiply-accumulate (MAC) operations required for sampling in the RBM during learning is compared to the number of synaptic operations (SynOps) in the spiking SSM during the MNIST learning task. At this task, the SSM requires fewer operations to reach the same accuracy during learning. Other necessary operations for the RBM, *e.g.* additions, random number generations, sigmoid function calls and weight updates were not taken into account here, and would further favor the SSM. One reason for the SynOp-MAC parity is that learning in the SSM requires fewer repetitions of data samples to reach the same accuracy compared to the RBM. Another reason is that only active connections in the SSM incur a SynOp. SynOp-MAC parity in the SSM is very promising for neuromorphic hardware implementations because a SynOp in dedicated hardware potentially consumes much less power than a MAC in a general purpose digital processor. Note that the discrete-time SSM is not competitive on this measure because it requires N^2 random number generations per Gibbs sampling step in addition to the same number of MACs as the RBM.

a stochastic component to Hopfield networks, but rather than the units themselves being noisy, the connections are noisy.

The stochastic synapse model we considered is blank-out noise, where the synaptic weight is multiplied by binary random variable. This is to be contrasted with additive noise, where the stochastic term such as a white noise process is added to the membrane potential. In artificial neural networks, multiplicative noise forces weights to become either sparse or invariant to rescaling [55]. When combined with rectified linear activation functions, multiplicative noise tends to increase sparsity [64]. Multiplicative noise makes neural responses sparser: In the absence of additive noise, neurons have activation functions very close to being rectified linear [67]. In the absence of a depolarizing input, the neuron cannot fire, and thus its synapses cannot potentiate. Consequently, if the parameters are initiated properly, many neurons will remain silent after learning.

Event-driven CD was first demonstrated in a spiking neural network that instantiated a Markov Chain Monte Carlo (MCMC) neural sampling of a target Boltzmann distribution [56]. The classification performance of the original model was limited by two properties. First, every neuron was injected with a (additive) noisy current of very large amplitude. Neurally speaking, this corresponded to a background Poisson spike train of 1000 Hz, which considerably increased the network activity in the system. Second, in spite of the large input noise, neurons fired periodically at large firing rates due to an absolute refractory period. This periodic firing evoked strong spike-to-spike correlations (synchrony) that were detrimental to the learning and the sampling. Consequently, the performance of eCD in an MNIST task was significantly lower than when standard CD was used (8.1% error rate).

The performance of the SSM in the MNIST hand-written digits task vastly improved over our previous results (4.4% error rate). The improvement in performance over our earlier results in [56] stems from at least two reasons: 1) Spike-to-spike decorrelations caused by the synaptic noise better condition the plasticity rule by preventing pair-wise synchonization; 2) Regularization, which mitigates overfitting and the co-adaptation of the units. In the machine learning community, blank-out noise is also known as DropConnect [69]. It was demonstrated to perform regularization, which helped achieve state-of-the-art results on several image recognition benchmarks. Note that the SSM model did not include domain-specific knowledge, which suggests that its performance may generalize to other problems and datasets.

4.1 Synaptic Unreliability in the Brain

The probabilistic nature of synaptic quantal release is a well known phenomenon [46]. Unreliability is caused by the probabilistic release of neurotransmitters at the pre-synaptic terminals [7]. Detailed slice and *in vivo* studies found that synaptic vesicle release in the brain can be extremely unreliable - typically 50% transmission rate and possibly as low as 10% in *in vivo* - at a given synapse [16, 15]. Such synaptic unreliability can be a major source of noise in the brain [19, 28, 72, 1]. Interestingly, one consequence of probabilistic synapses is that, via recurrent interactions in the networks, synaptic unreliability would be the cause of pre-synaptically caused noise.

In the SSM, the multiplicative effect of the blank-out noise is manifested in the pre-synaptic input by making its variance dependent on the synaptic weights and the network states. Our theoretical analysis suggests that, in the SSM's regime of operation, increased variability has the effect of reducing the sensitivity of the neuron to other synaptic inputs by flattening the neural transfer curve. With multiplicative noise, input variability can be high when pre-synaptic neurons with strong synaptic weights are active. In the SSM, such an activity pattern emerges when the probability of a given state under the learned model and the sensory data is high. That case suggests that the network has reached a good estimate and should not be easily modified by other evidence,

which is the case when the neural transfer curve is flatter. Synaptic unreliability can thus play the role of a dynamic normalization mechanism in the neuron with direct implications on probabilistic inference and action selection in the brain.

Aithchison and Latham suggested the "synaptic sampling hypothesis" whereby pre-synaptic spikes would draw samples from a distribution of synaptic weights [2]. This process could be a mechanism used in the brain to represent uncertainty over the parameters of a learned model [4]. Stochasticity can be carried to the learning dynamics as well. Recent studies point to that fact, when learning the blank-out probability at the synapses is also learned [6, 5], the learned neural generative models capture richer representations, especially when labeled data is sparse. Kappel and colleagues also showed that stochasticity in the learning dynamics improves generalization capability of neural network.

The blank-out noise model used in this work is a particular case of the studies above, whereby the weights of the synapses are either zero or the value of the stored weight with a fixed probability. In contrast to previous work, we studied the learning dynamics under this probabilistic synapse model within an otherwise deterministic recurrent neural network. Besides the remarkable fact that synaptic stochasticity alone is sufficient for sampling, it enables robust learning of sparse representations and an efficient implementation in hardware.

Related Work on Mapping Machine Learned Models onto Neuromorphic Hardware Many approaches for configuring spiking neural networks rely on mapping pre-trained artificial neural networks onto spiking neural networks using a firing rate code [57, 26, 20, 42]. Many recent work show that the approximations incurred in neuromorphic platforms [50, 56, 20, 26, 57, 23, 49, 49] including reduced bit precision [66, 53, 49, 49] have a minor impact on performance. Standard artificial neural networks such as deep convolutional neural networks and recurrent neural networks trained with Rectified linear units (ReLU) can be mapped on spiking neurons by exploiting the threshold-linear of integrate & fire neurons [26, 20]. Such mapping techniques have the advantage that they can leverage the capabilities of existing machine learning frameworks such as Caffe [44] or pylearn2 [35] for neuromorphic platforms. Although mapping techniques do not offer a solution for on-line, real-time learning, they resulted in the best performing spike-based implementations on standard machine learning benchmarks such as MNIST and CIFAR.

Related work on Online Learning with Spiking Neurons Training neural networks is a very time- and energy-consuming operation, often requiring multiple days on a computing cluster to produce state-of-the-art results. Using neuromorphic architectures for learning neural networks can have significant advantages from the perspectives of scalability, power dissipation and real-time interfacing with the environment. While embedded synapses require additional chip resources and usually prohibits the implementation of more complex rules, a recent survey of software simulators argues that dedicated learning hardware is a prerequisite for real-time learning (or faster) in spiking networks [73]. This is because the speed-up margin of parallelism encounters a hard boundary due to latencies in the inter-process communications.

The SSM is an ideal candidate for *embedded*, online learning, where plasticity is implemented locally using a dedicated on-chip device [9]. These operate on local memory (*e.g.* synaptic weights) using local information (*e.g.* neural events) which will lead to more scalable, faster and more power efficient learning compared to computer clusters, and the ability to adapt in real-time to changing environments.

Previous work addressed on-line learning with spiking neurons. A hierarchical spiking neural network with a STDP-like learning rule achieved 8% error rate on the MNIST task [14]. Diehl and

Cook demonstrated the best results on unsupervised learning with spiking neural networks so far [25]. Their network model is comparable to competitive learning algorithms where each neuron learns a representation of a portion of the input space. Their architecture could achieve up to 5% error rate. The SSM outperformed this best result using a much smaller number of neurons and synapses and relying on dynamics that are more amenable to hardware implementations. However, the number of repetitions to reach this performance using the SSM was larger than the above studies (512,000 presentations vs 40,000 in [25]). Our neural sampling based architecture with stochastic neurons and deterministic synapses achieved peak results after 15,000 samples [56], suggesting that the slowness is caused by the stochastic connections. Similar results have been observed using the DropConnect algorithm [69].

4.2 Implementations of Synaptic Unreliability in Neuromorphic Hardware

At least four studies reported the implementation of blank-out synapses for neuromorphic systems using the Address Event Representation (AER) [34, 22, 21, 50]. In these studies, synaptic unreliability was mainly used as a mechanism for increasing the resolution of the synaptic weights in hardware (which is often binary). In fact, the mean of a synaptic current produced by a stochastic synapse is the probability times the weight of the synapse. By allowing the probability to take values with high precision, the effective resolution of the synapse weight can be increased. The downside is that this approach is valid only when the neural computations are rate-based, such as in the neural engineering framework [27] where synaptic unreliability in neuromorphic systems was primarily applied [22, 21].

In rate-based models, the variability introduced by stochastic synapses is dealt with by averaging over large populations of neurons or by taking temporal averages. Implementations based on firing rate codes thus disregard spike times. From a hardware perspective, firing rate codes often raise the question whether a spike-based neuromorphic platform is justifiable over a direct, dedicated implementation of the machine learning operations, or even a dedicated implementation of the rate dynamics [70]. In contrast, codes based on neural sampling, synaptic sampling or phase critically depend on spike statistics or the precise timing of the spikes. For example, in the SSM, synaptic unreliability and the variability that it causes are an integral part of the sampling process. The variability introduced by the multiplicative property of synaptic noise is exploited both as a mechanism that generates sigmoidal activation and that improves learning. Results from the network dynamics suggest that the introduced variability generates sparser representations, and in some cases are insensitive to parameter rescaling. Thus, our work suggests that synaptic unreliability can play much more active roles in information processing in neuromorphic hardware systems.

σ	Noise amplitude	spiking SSM, visible neurons	4.47 nA
p	Blank-out probability at synapse	all models	.5
τ_r	Refractory period	all spiking SSM	4 ms
τ_{syn}	Time constant of recurrent, and bias synapses.	all spiking SSM	4 ms
τ_{br}	"Burn-in" time of the neural sampling	all spiking SSM	10 ms
g_L	Leak conductance	all spiking SSM	1 nS
u_{rst}	Reset Potential	all spiking SSM	0 V
C	Membrane capacitance	all spiking SSM	1 pF
θ	Firing threshold	all spiking SSM	100 mV
$2T$	Epoch duration	all spiking SSM	100 ms
W	Initial weight matrix	all spiking SSM	$N(0, .3)$
		all dSSM, RBM	$U(0, .1)$
b_v, b_h	Initial bias for layer v and h	all spiking SSM	$-.15$
T_{sim}	Simulation time per epoch	all spiking SSM	100 s
N_v, N_h	Number of visible and hidden units	all models,	$794, 500$
		except in Fig. 1	$5, 5$
N_c	Number of class label units	all models	10
$N_{samples}$	Total number of MNIST sample presentations	Fig. 4, spiking SSM	512000
		Fig. 4, dSSM, RBM	$75 \cdot 10^6$
τ_{STDP}	Learning time window	all models	10 ms
ϵ_q	Learning rate for W	all spiking SSM	$3.85 \cdot 10^{-6}$
ϵ		all dSSM, RBM	.025
ϵ_b	Learning rate for b_v, b_h	all spiking SSM	$1.43 \cdot 10^{-5}$
		all dSSM, RBM	.025
n_{batch}	Batch size	all dSSM, RBM	50
		spiking SSM	(1)
W	Distribution of weight parameters	Fig. 1	$N(-.3, 1.5)$
b_h, b_v	Distribution of bias parameters	Fig. 1	$N(0, 1.5)$

Table 2: (1) There is no possibility of batch learning in the spiking SSM

5 Acknowledgments

We thank Friedemenn Zenke for support on the Auryn simulator and discussion.
This work funded by the National Science Foundation (NSF CCF-1317373, EN, BP, SJ, GC), the Office of Naval Research (ONR MURI 14-13-1-0205, EN, BP), and Intel Corporation (EN, GC).

References

[1] L.F. Abbott and W.G. Regehr. Synaptic computation. *Nature*, 431:796–803, October 2004.

[2] L. Aitchison and P. Latham. The synaptic sampling hypothesis. *Cosyne Abstracts*, 2013.

[3] L. Aitchison and P. E. Latham. Synaptic sampling: A connection between PSP variability and uncertainty explains neurophysiological observations. *ArXiv e-prints*, May 2015.

[4] Laurence Aitchison and Peter E Latham. Bayesian synaptic plasticity makes predictions about plasticity experiments in vivo. *arXiv preprint arXiv:1410.1029*, 2014.

[5] M Al-Shedivat*, R. Naous, E. Neftci, G. Cauwenberghs, and K.N. Salama. Inherently stochastic spiking neurons for probabilistic neural computation. In *IEEE EMBS Conference on Neural Engineering*, Apr 2015. (accepted).

[6] M. Al-Shedivat, E. Neftci, and G. Cauwenberghs. Neural generative models with stochastic synapses capture richer representations. In *Cosyne Abstracts*, May 2015. (accepted).

[7] Christina Allen and Charles F Stevens. An evaluation of causes for unreliability of synaptic transmission. *Proceedings of the National Academy of Sciences*, 91(22):10380–10383, 1994.

[8] D.J. Amit and N. Brunel. Dynamics of a recurrent network of spiking neurons before and following learning. *Network: Computation in Neural Systems*, 8(4):373–404, 1997.

[9] R.M. Azghadi, N. Iannella, S.F. Al-Sarawi, G. Indiveri, and D. Abbott. Spike-based synaptic plasticity in silicon: design, implementation, application, and challenges. *Proceedings of the IEEE*, 102(5):717–737, 2014.

[10] Thomas M Bartol, Cailey Bromer, Justin P Kinney, Michael A Chirillo, Jennifer N Bourne, Kristen M Harris, and Terrence J Sejnowski. Hippocampal spine head sizes are highly precise. *bioRxiv*, page 016329, 2015.

[11] Ben Varkey Benjamin, Peiran Gao, Emmett McQuinn, Swadesh Choudhary, Anand R Chandrasekaran, J Bussat, Rodrigo Alvarez-Icaza, John V Arthur, PA Merolla, and Kwabena Boahen. Neurogrid: A mixed-analog-digital multichip system for large-scale neural simulations. *Proceedings of the IEEE*, 102(5):699–716, 2014.

[12] J. Bergstra, O. Breuleux, F. Bastien, P. Lamblin, R. Pascanu, G. Desjardins, J. Turian, D. Warde-Farley, and Y. Bengio. Theano: a CPU and GPU math expression compiler. In *Proceedings of the Python for Scientific Computing Conference (SciPy)*, volume 4, 2010.

[13] Pietro Berkes, Gergő Orbán, Máté Lengyel, and József Fiser. Spontaneous cortical activity reveals hallmarks of an optimal internal model of the environment. *Science*, 331(6013):83–87, 2011.

[14] Michael Beyeler, Nikil D. Dutt, and Jeffrey L. Krichmar. Categorization and decision-making in a neurobiologically plausible spiking network using a stdp-like learning rule. *Neural Networks*, 48:109–124, 2013.

[15] J Gerard G Borst. The low synaptic release probability in vivo. *Trends in neurosciences*, 33(6):259–266, 2010.

[16] Tiago Branco and Kevin Staras. The probability of neurotransmitter release: variability and feedback control at single synapses. *Nature Reviews Neuroscience*, 10(5):373–383, 2009.

[17] N. Brunel. Dynamics of sparsely connected networks of excitatory and inhibitory spiking neurons. *Journal of computational neuroscience*, 8(3):183–208, 2000.

[18] N. Brunel and V. Hakim. Fast global oscillations in networks of integrate-and-fire neurons with low firing rates. *Neural Computation*, 11(7):1621–1671, 1999.

[19] William H Calvin and CHARLES F Stevens. Synaptic noise and other sources of randomness in motoneuron interspike intervals. *J Neurophysiol*, 31(4):574–587, 1968.

[20] Yongqiang Cao, Yang Chen, and Deepak Khosla. Spiking deep convolutional neural networks for energy-efficient object recognition. *International Journal of Computer Vision*, pages 1–13, 2014.

[21] S. Choudhary, S. Sloan, S. Fok, A. Neckar, E. Trautmann, P. Gao, T. Stewart, C. Eliasmith, and K. Boahen. Silicon neurons that compute. In A. Villa, W. Duch, P. Érdi, F. Masulli, and G. Palm, editors, *Artificial Neural Networks and Machine Learning – ICANN 2012*, volume 7552 of *Lecture Notes in Computer Science*, pages 121–128. Springer Berlin / Heidelberg, 2012.

[22] Federico Corradi, Chris Eliasmith, and Giacomo Indiveri. Mapping arbitrary mathematical functions and dynamical systems to neuromorphic vlsi circuits for spike-based neural computation. In *Circuits and Systems (ISCAS), 2014 IEEE International Symposium on*, pages 269–272. IEEE, 2014.

[23] Srinjoy Das, Bruno Umbria Pedroni, Paul Merolla, John Arthur, Andrew S Cassidy, Bryan L Jackson, Dharmendra Modha, Gert Cauwenberghs, and Ken Kreutz-Delgado. Gibbs sampling with low-power spiking digital neurons. *arXiv preprint arXiv:1503.07793*, 2015.

[24] G. Deco, V.K. Jirsa, P.A. Robinson, M. Breakspear, and K. Friston. The dynamic brain: From spiking neurons to neural masses and cortical fields. *PLoS Comput Biol*, 4(8):e1000092, 2008.

[25] Peter U Diehl and Matthew Cook. Unsupervised learning of digit recognition using spike-timing-dependent plasticity. *Frontiers in Computational Neuroscience*, 9:99, 2015.

[26] Peter U Diehl, Daniel Neil, Jonathan Binas, Matthew Cook, Shih-Chii Liu, and Michael Pfeiffer. Fast-classifying, high-accuracy spiking deep networks through weight and threshold balancing. In *International Joint Conference on Neural Networks (IJCNN),*, pages 1–8. IEEE, 2015.

[27] C. Eliasmith and C.H. Anderson. *Neural engineering: Computation, representation, and dynamics in neurobiological systems.* MIT Press, 2004.

[28] A Aldo Faisal, Luc PJ Selen, and Daniel M Wolpert. Noise in the nervous system. *Nature Reviews Neuroscience*, 9(4):292–303, 2008.

[29] József Fiser, Pietro Berkes, Gergő Orbán, and Máté Lengyel. Statistically optimal perception and learning: from behavior to neural representations. *Trends in cognitive sciences*, 14(3):119–130, 2010.

[30] S. Fusi and M. Mattia. Collective behavior of networks with linear (VLSI) integrate and fire neurons. *Neural Computation*, 11:633–52, 1999.

[31] Crispin W Gardiner. Handbook of stochastic methods. 2012.

[32] W. Gerstner and W. Kistler. *Spiking Neuron Models. Single Neurons, Populations, Plasticity.* Cambridge University Press, 2002.

[33] M. Giulioni, F. Corradi, V. Dante, and P. del Giudice. Real time unsupervised learning of visual stimuli in neuromorphic vlsi systems. *Scientific Reports*, 5:14730 EP –, Oct 2015. Article.

[34] D.H. Goldberg, G. Cauwenberghs, and A.G. Andreou. Probabilistic synaptic weighting in a reconfigurable network of VLSI integrate-and-fire neurons. *Neural Networks*, 14(6–7):781–793, Sep 2001.

[35] Ian J. Goodfellow, David Warde-Farley, Pascal Lamblin, Vincent Dumoulin, Mehdi Mirza, Razvan Pascanu, James Bergstra, Frédéric Bastien, and Yoshua Bengio. Pylearn2: a machine learning research library. *arXiv preprint arXiv:1308.4214*, 2013.

[36] Simon S Haykin. Neural networks: A comprehensive foundation, 1999.

[37] G.E. Hinton, S. Osindero, and Y.W. Teh. A fast learning algorithm for deep belief nets. *Neural computation*, 18(7):1527–1554, 2006.

[38] G.E. Hinton and R.R. Salakhutdinov. Reducing the dimensionality of data with neural networks. *Science*, 313(5786):504–507, 2006.

[39] G.E. Hinton and T.J. Sejnowski. Learning and relearning in boltzmann machines. *MIT Press, Cambridge, Mass*, 1:282–317, 1986.

[40] Geoffrey E Hinton. Training products of experts by minimizing contrastive divergence. *Neural computation*, 14(8):1771–1800, 2002.

[41] J.J. Hopfield. Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the national academy of sciences*, 79(8):2554–2558, 1982.

[42] Eric Hunsberger and Chris Eliasmith. Spiking deep networks with lif neurons. *arXiv preprint arXiv:1510.08829*, 2015.

[43] G. Indiveri, B. Linares-Barranco, T.J. Hamilton, A. van Schaik, R. Etienne-Cummings, T. Delbruck, S.-C. Liu, P. Dudek, P. Häfliger, S. Renaud, J. Schemmel, G. Cauwenberghs, J. Arthur, K. Hynna, F. Folowosele, S. Saighi, T. Serrano-Gotarredona, J. Wijekoon, Y. Wang, and K. Boahen. Neuromorphic silicon neuron circuits. *Frontiers in Neuroscience*, 5:1–23, 2011.

[44] Yangqing Jia, Evan Shelhamer, Jeff Donahue, Sergey Karayev, Jonathan Long, Ross Girshick, Sergio Guadarrama, and Trevor Darrell. Caffe: Convolutional architecture for fast feature embedding. *arXiv preprint arXiv:1408.5093*, 2014.

[45] David Kappel, Stefan Habenschuss, Robert Legenstein, and Wolfgang Maass. Network plasticity as bayesian inference. *arXiv preprint arXiv:1504.05143*, 2015.

[46] B. Katz. *Nerve, muscle, and synapse*. McGraw-Hill New York, 1966.

[47] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.

[48] Honglak Lee, Chaitanya Ekanadham, and Andrew Y Ng. Sparse deep belief net model for visual area v2. In *Advances in neural information processing systems*, pages 873–880, 2008.

[49] Daniel Marti, Mattia Rigotti, Mingoo Seok, and Stefano Fusi. Energy-efficient neuromorphic classifiers. *arXiv preprint arXiv:1507.00235*, 2015.

[50] Paul A Merolla, John V Arthur, Rodrigo Alvarez-Icaza, Andrew S Cassidy, Jun Sawada, Filipp Akopyan, Bryan L Jackson, Nabil Imam, Chen Guo, Yutaka Nakamura, et al. A million spiking-neuron integrated circuit with a scalable communication network and interface. *Science*, 345(6197):668–673, 2014.

[51] S. Moradi and G. Indiveri. An event-based neural network architecture with an asynchronous programmable synaptic memory. *IEEE Transactions on Biomedical Circuits and Systems*, March 2013.

[52] R. Moreno-Bote. Poisson-like spiking in circuits with probabilistic synapses. *PLoS computational biology*, 10(7):e1003522, 2014.

[53] Lorenz K Muller and Giacomo Indiveri. Rounding methods for neural networks with low resolution synaptic weights. *arXiv preprint arXiv:1504.05767*, 2015.

[54] Vinod Nair and Geoffrey E. Hinton. 3d object recognition with deep belief nets. In Y. Bengio, D. Schuurmans, J.D. Lafferty, C.K.I. Williams, and A. Culotta, editors, *Advances in Neural Information Processing Systems 22*, pages 1339–1347. Curran Associates, Inc., 2009.

[55] Eric Nalisnick, Anima Anandkumar, and Padhraic Smyth. A scale mixture perspective of multiplicative noise in neural networks. *arXiv preprint arXiv:1506.03208*, 2015.

[56] E. Neftci, S. Das, B. Pedroni, K. Kreutz-Delgado, and G. Cauwenberghs. Event-driven contrastive divergence for spiking neuromorphic systems. *Frontiers in Neuroscience*, 7(272), Jan. 2014.

[57] P O'Connor, D. Neil, S.-C. Liu, T. Delbruck, and M. Pfeiffer. Real-time classification and sensor fusion with a spiking deep belief network. *Frontiers in Neuroscience*, 7(178), 2013.

[58] J. Park, S. Ha, T. Yu, E. Neftci, and G. Cauwenberghs. A 65k-neuron 73-mevents/s 22-pj/event asynchronous micro-pipelined integrate-and-fire array transceiver. In *Biomedical Circuits and Systems Conference (BioCAS)*. IEEE, Oct. 2014.

[59] Emanuel Parzen. *Stochastic processes*, volume 24. SIAM, 1999.

[60] Mihai A Petrovici, Johannes Bill, Ilja Bytschok, Johannes Schemmel, and Karlheinz Meier. Stochastic inference with deterministic spiking neurons. *arXiv preprint arXiv:1311.3211*, Nov 2013.

[61] T. Pfeil, T. C. Potjans, S. Schrader, W. Potjans, J. Schemmel, M. Diesmann, and K. Meier. Is a 4-bit synaptic weight resolution enough? - constraints on enabling spike-timing dependent plasticity in neuromorphic hardware. *Frontiers in Neuroscience*, 6, 2012.

[62] Dimitri Probst, Mihai A Petrovici, Ilja Bytschok, Johannes Bill, Dejan Pecevski, Johannes Schemmel, and Karlheinz Meier. Probabilistic inference in discrete spaces can be implemented into networks of lif neurons. *Frontiers in computational neuroscience*, 9, 2015.

[63] A. Renart, N. Brunel, and X. Wang. *Computational Neuroscience: A Comprehensive Approach*, chapter Mean field theory of irregularly spiking neuronal populations and working memory in recurrent cortical networks, pages 431–490. Chapman and Hall, Boca Raton, 2003.

[64] Jan Rudy, Weiguang Ding, Daniel Jiwoong Im, and Graham W Taylor. Neural network regularization via robust weight factorization. *arXiv preprint arXiv:1412.6630*, 2014.

[65] J. Schemmel, D. Brüderle, A. Grübl, M. Hock, K. Meier, and S. Millner. A wafer-scale neuromorphic hardware system for large-scale neural modeling. In *International Symposium on Circuits and Systems, ISCAS 2010*, pages 1947–1950. IEEE, 2010.

[66] Evangelos Stromatias, Daniel Neil, Michael Pfeiffer, Francesco Galluppi, Steve B Furber, and Shih-Chii Liu. Robustness of spiking deep belief networks to noise and reduced bit precision of neuro-inspired hardware platforms. *Frontiers in neuroscience*, 9, 2015.

[67] Henry C Tuckwell. *Introduction to theoretical neurobiology: volume 2, nonlinear and stochastic theories*, volume 8. Cambridge University Press, 2005.

[68] C. van Vreeswijk and H. Sompolinsky. Chaos in neuronal networks with balanced excitatory and inhibitory activity. *Science*, 274(5293):1724–1726, December 1996.

[69] Li Wan, Matthew Zeiler, Sixin Zhang, Yann L Cun, and Rob Fergus. Regularization of neural networks using dropconnect. In *Proceedings of the 30th International Conference on Machine Learning (ICML-13)*, pages 1058–1066, 2013.

[70] Runchun Wang, Chetan Singh Thakur, Tara Julia Hamilton, Jonathan Tapson, and André van Schaik. A neuromorphic hardware architecture using the neural engineering framework for pattern recognition. *arXiv preprint arXiv:1507.05695*, 2015.

[71] X. Wang. Synaptic basis of cortical persistent activity: the importance of NMDA receptors to working memory. *J. Neurosci.*, 19:9587–9603, November 1999.

[72] Yosef Yarom and Jorn Hounsgaard. Voltage fluctuations in neurons: signal or noise? *Physiological Reviews*, 91(3):917–929, 2011.

[73] F. Zenke and W. Gerstner. Limits to high-speed simulations of spiking neural networks using general-purpose computers. *Frontiers in neuroinformatics*, 8, 2014.